
Third-party information provided to you courtesy of Dell.



DELL/EMC CX SERIES INITIAL CONFIGURATION

BEST PRACTICES PLANNING

Abstract

This white paper provides basic best practice recommendations for Dell/EMC CX Series® storage system settings. The recommendations provide parameters that will apply to 80 percent of the configurations installed.

February 2007

Copyright © 2005, 2007 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

CONTENTS

Executive summary	5
Introduction	6
Audience	6
General hardware	6
Drive type differentiation	6
RAID groups	7
RAID group general	7
RAID group settings	8
LUNs	9
LUN creation settings	9
Special LUN considerations	10
LUN expansion: MetaLUNs	10
RAID type to file type match	11
Application profiles	11
Drive count	13
Drive types	13
Hot spares	13
Drive type	13
Vault and configuration drives	14
Simple and efficient storage system configurations	14
Storage Groups	17
Create a Storage Group	17
Storage Group settings	17
Caching	18
Cache settings and other SP information	18
Software online upgrades	18
Monitoring	19
Management system	19
Off-array management server	19
In-array management	20
Host management	20
Host Agent settings	20
Management scope	20
Local diagnosis EMCRemote	20
Event notification	20
Sessions limit	23
Security	24
Storage systems	24
Management	24
Domains	24
Access levels	24
Audit logging	24
Password control	24
Access lost or forgotten passwords	24
Connectivity	25
Management	25
Protocols	25

iSCSI connections.....	25
NIC cards.....	25
TOE cards.....	25
iSCSI security.....	25
Fibre Channel connections.....	25
Emulex HBAs.....	26
QLogic HBAs: QLogic settings.....	26
Data path.....	26
Host bus adapters.....	26
Dell/EMC CX Series SP ports.....	26
Zoning.....	27
Port-to-HBA relationships.....	27
Fabric connectivity.....	27
Multipath software.....	28
LUN presentation.....	28
RAID group to back-end bus connections.....	28
Back-end connectivity.....	29
Clustering.....	29
Sun clusters.....	29
VERITAS clusters.....	29
IBM clusters.....	29
Windows clusters: Microsoft Cluster Server	30
Linux (RedHat) clusters.....	30
HP TruCluster.....	30
Scaling: Storage system sizes.....	30
Small storage system.....	30
Medium storage system.....	30
Large storage system.....	30
Host connectivity.....	30
Exchange environments.....	31
Disk-array enclosure (disk shelves)	31
Availability.....	31
Drive types per DAE.....	31
Drives per DAE.....	31
Drive types per storage system.....	31
Drives per RAID group.....	32
Cabinets.....	32
Conclusion.....	32
References.....	32
Replication software.....	32
Other references.....	32

EXECUTIVE SUMMARY

This white paper is provided to expedite deployment of the Dell/EMC® CX3 Series. It defines how to best implement the new/enhanced technology into customer environments. The paper provides best practices that will satisfy 80 percent of customer configurations; they meet the *80/20 rule*. (The 80/20 rule states if X is done, X will satisfy 80 percent of the situations or configurations.)

The guidelines in this white paper are not rules but rather recommendations within the rules. Most implementations should be successful if these recommendations are followed. Using this paper, EMC can partner with customers to apply the customer's own best practices policies (if any). This is a "living document," changing with experience and the addition of new features.

This paper is broken down into five sections:

- General Hardware
- Monitoring
- Security
- Connectivity
- Scaling

There may be some overlap between sections as some capabilities may overlap.

The paper also notes cases where a configuration or feature may fall outside the 80 percent rule. Most often, these cases are based on performance or configuration limitations.

Best practice settings for performance are *not* included in this document. Although they are taken into consideration here, they are not the primary focus. The primary focus of this paper is function and ease of installation—not fine-tuning. To access the best practices for performance, refer to *EMC Dell/EMC CX Series Best Practices for Fibre Channel Storage* on EMC® Powerlink®.

INTRODUCTION

Customers are requiring simpler installs on Dell/EMC CX Series CX3 Series storage systems. One way to accommodate this requirement, without reducing the power and flexibility of Dell/EMC CX Series storage systems, is to provide best practices that meet the 80/20 rule. Many customers have requested that EMC provide basic recommendations for configurations, and have suggested that EMC “put a stake in the ground” for most configurations. This paper accomplishes that for most configurations, from an operational and management perspective.

The primary environments this paper addresses are: file and print, basic database online transaction processing (OLTP), and data warehouses. Video environments’ special modifications will be noted as such when needed.

The settings included in this document are not rules but rather guidelines within rules. For example:

- **Rule** – RAID 5 supports up to 16 disks.
- **Guideline** – Set up RAID 5 in groups of five or nine disks.

If you want to see the rules, refer to the *EMC Dell/EMC CX Series Open System Configuration Guide* (either the documentation or the *EMC Support Matrix*).

Deviations from the best practice are allowed and encouraged, especially for experienced users. However, if you want a simple guide to maximize a successful installation, follow the best practice. The best practices are instructions, such as “set this up” or “use that setting.” References to best practices point to other documents that were written to address the specific needs of a particular application or to maximize performance of the Dell/EMC CX Series.

Tables include recommended parameter settings for particular operations, such as creating a RAID group. The tables include the attribute name (as seen in Navisphere® User Interface), the best practice setting, the default setting and any notes relating to the parameter.

Audience

This white paper is intended for EMC field personnel and customers who want to configure a Dell/EMC CX Series for general use.

GENERAL HARDWARE

The Dell/EMC CX Series is a very flexible storage system that allows the administrator to configure it many different ways.

Drive type differentiation

For the purposes of configuration, the following characteristics can be used to differentiate between drives when identifying different configurations. Different models have different drive types, sizes, and rotational speeds.

- **Technology** - EMC supports two different drive types within the Dell/EMC CX Series storage system: Fibre Channel (FC) and Advanced Technology Attachment (ATA). Both can be contained within one Dell/EMC CX Series but only one type can be in any one DAE.
- **Size** - EMC provides different sizes for drives based on type. EMC Fibre Channel drive sizes are: 73 GB, 146 GB, 300 GB, and 500 GB. ATA/SATA-II drives are available in 500 GB capacity. Larger size drives of the same type are allowed as replacement drives in the same RAID group. Remember that replacing a drive of one size with one that is larger only provides you with the capacity of the smallest

capacity drive in the RAID group. This usually coincides with the capacity of the replaced drive.

- **Rotational speed** – ATA and SATA-II drives currently have a rotational speed of 7,200 revolutions per minute (rpm). Fibre Channel drives that rotate at 7,200, 10,000, or 15,000 rpm are available.
- **Backend speed** – EMC provides two different back-end speed drives for CX3 systems. The two available speeds are 4 Gb/s or 2 Gb/s.
- **Vault and database drives** – Drives 0, 1, 2, 3 and 4 in the first DAE in a Dell/EMC CX Series contain the vault. These are 10,000 or 15,000 rpm FC drives and can be of any capacity or back-end speed. The vault drives are special because they contain configuration information, the write cache dump area, and other information that is used by the Dell/EMC CX Series. The vault drives are also the storage area that protects the write cache in case of a power failure.

RAID groups

This section provides specific guidance for RAID groups.

RAID group general

1. Put the same capacity drives in the same RAID group. The RAID group will be created at the capacity of the lowest capacity drive. Extra space on the bigger drives will be wasted.
2. Put same speed drives in the RAID group. This maximizes the service capability of the group.
3. Mixing ATA and Fibre Channel drives is not allowed in the same RAID group. This is enforced by the Navisphere management software.

Fibre Channel drives

For 80 percent of the general-usage configurations use RAID 5: 4+1, 8+1. Any number of drives between 3 and 16 can be used, but RAID groups of this size are generally easier to configure and manage within a single DAE chassis.

For 80 percent of OLTP configurations, use RAID 5: 4+1, 8+1 for data files (index and/or tables) and RAID 1/0: 1+1 for OLTP log files. Refer to “Application profiles” for application-specific modifications to these recommendations.

ATA drives

For ATA and SATA-II drives use RAID 3 4+1 or 8+1. ATA and SATA-II drives should be used for archival or media applications, whose I/O patterns are sequential in nature. FLARE® has been tuned to maximize this capability for ATA and SATA-II drives when using RAID 3. When using RAID 3, ensure that write cache is enabled for the LUNs.

All LUNs created from ATA drives in the same DAE should be assigned the same default storage processor (SP). This guideline does not apply to SATA-II drives.

RAID group settings

This type of table is used throughout the document. It follows what Navisphere Manager presets for a particular configuration area. Table headings indicate the following:

- *Set* is the 80 percent recommendation for the particular parameter.
- *Default Setting* is what the system presents to the user.
- *Notes* provide either extra explanation or mitigating circumstances that could cause different settings.

Table 1. RAID group settings

	Set	Default Setting	Notes
Create a RAID group			
Group ID	Default	Next lowest available number	
Number of disks	9	1	This allows an 8+1 RAID 5.
RAID type	Default	Unbound	
Automatically destroy When last LUN unbound	Default	Field not checked	
Expansion Defragmentation priority	Default	Medium	
Disk selection	Default	Automatic	The Dell/EMC CX Series will choose a group of disks that fits the requirements. Use manual selection and select the desired disks if you are in an environment with mixed disk sizes. If you want to use manual, go to "Simple and efficient storage system configurations" for guidance.

LUNs

LUNs are created from a RAID group. The RAID type for a RAID group will be assigned when the first LUN is bound on the RAID group. LUNs can be bound in many different ways on a Dell/EMC CX Series. For maximum availability all LUNs should be bound with protected RAID, such as RAID 1/0, or RAID 5. A number of hot spares should also be bound (see hot spares section).

LUN creation settings

Table 2 lists the recommended settings for creating LUNs.

Table 2. LUN creation settings

	Set	Default setting	Notes
Bind a LUN			
RAID type	RAID 5, RAID 3, RAID 1 or RAID 1/0	RAID 5	Refer to the "Application" section. RAID 5 is considered the best cost/performance balance for protected RAID.
RAID Group ID	Choose any	1	Choose the RAID Group ID you wish to put this LUN into. If none exists, click on new .
LUN ID	See note	Next lowest available	Try to use an even and odd numbering system to reflect the default owner: For example, use even for SPA and odd for SP B.
No Initial Verify	Default	No (box is blank)	Do not send data to the LUN until the background verify operation is complete.
Rebuild priority	Default	ASAP	
Verify priority	Default	ASAP	
Default owner	Auto	Either SP A or SP B depending on the LUN allocation to each	
Alignment offset	See note	0	The alignment offset is operating system (OS) dependent. Windows: 63 For other OSs or use with replication software, use the default and see the section directly below this table for specific recommendations.
Read and write cache	Default	Enabled	
Enable auto assign	Default	No (box blank)	
LUN size	See note	1 GB	LUN size must be compatible with the operating system.
Number of LUNs to bind	See note	1	When adding LUNs to an existing system. select 1 .

For Windows systems:

If you plan to use replication software (SnapView, SAN Copy™ or MirrorView™) with a LUN, use the default offset and align using diskpar. The white paper *EMC Dell/EMC CX Series Best Practices for Fibre Channel Storage* has more details.

Special LUN considerations

When binding LUNs destined for the reserved LUN pool, give these LUNs high numbers. Place these LUNs in the “private LUN number” range. If the array to be configured supports a max of 2,048 addressable LUNs, then these reserved LUNs should start at 2100 and increase as more reserved LUNs are added.

LUN expansion: MetaLUNs

Create metaLUNs from LUNs in separate RAID 5 RAID groups. Use the striping method and wait for the extra space. This will provide the best performance in the majority of environments.

Striping metaLUNs typically provides high performance but requires the stripe to be built before the extra space becomes available. Concatenation provides immediate space availability but will not balance the I/Os across all the drives. The extra space must be used to get the added disk benefit from the concatenation.

The two can be combined in a concatenated stripe. This provides both immediate space availability and the potential performance benefit of striping.

For the stripe multiplier use:

- For four-disk LUNs use 4 – a four-disk LUN is a 4+1 RAID 3 or RAID 5 or a 4+4 RAID 1/0.
- For eight-disk LUNs use 2 – an eight-disk LUN is an 8+1 RAID 3 or RAID 5.
- If RAID 1, set to **2**.
- A good general purpose number is **4**.

LUNs should be expanded onto “like” RAID groups (that is, expand a 4+1 R5 LUN into other 4+1 R5 RAID groups to retain similar geometry).

For user capacity, use the current capacity for planned metaLUN creation. Use the maximum capacity setting for unplanned metaLUN creation. While metaLUN expansion is supported on all Dell/EMC CX Series arrays, each operating system has separate restrictions regarding the support of online capacity expansion. Refer to your operating system vendor for details.

All member LUNs of a metaLUN should be assigned to the same default SP.

MetaLUN settings

Table 3 describes the user-defined and default parameters for metaLUNs

Table 3. MetaLUN settings

	Set	Default setting	Notes
LUN Expansion			
Type of expansion	Striping	Striping	The user will need to wait for the striping to complete before the space can be made available to the host.
Select LUN			All LUNs that meet the rules will be displayed. Select one LUN.
User capacity	GB	GB	This part of the capacity represents the units of measure.
User capacity (size)	Maximum capacity	Maximum capacity	
MetaLUN name	Follow customer naming practices	Same as base LUN name	Can be changed during the creation of the metaLUN
Default owner	Current SP	Current SP	The default will be the same as the current owner of the base LUN.
Expansion rate	ASAP	Default	
Element multiplier	default	4	This sets the largest I/O a single LUN can receive within a metaLUN
Alignment offset	default	Current setting	See section on creating LUNs for description.
Enable Auto Assign	default	Box unchecked	

RAID type to file type match

Most customer applications should be using RAID 5. This provides the best reliability and performance at the most cost-effective level. Variations from RAID 5 may occur based on the application profile.

Application profiles

EMC strongly recommends protected RAID for all applications.

There are four main factors that affect the configuration:

- **Capacity** – How much space is required?
 - Small environments: up to 4 TB
 - Medium environments: 4 – 11 TB
 - Large environments: 11+TBYour view of these ratings may vary, but the basis for these choices is:
 - A small configuration is one to two DAEs.
 - A medium configuration is two to eight DAEs with ATA drives.
 - A large configuration covers the rest.
- **Throughput** – How many I/Os per second (IOPS) will the load generate? Small I/O sizes can generally generate more IOPS than large I/O sizes.
- **Bandwidth** – How much data—in megabytes per second (MB/s)—will the load generate? Small I/O sizes tend not to move as much data as large I/O sizes.

-
- **Response Time** – This is the time it takes to complete a transaction. A transaction can be one or many I/Os. Typical response time is calculated on a per-I/O basis. Online transaction processing (OLTP) environments typically measure response time by a full transaction, which is usually many I/Os.

Online transaction processing (OLTP)

OLTP environments have special behaviors and requirements. Typically, a database will have two primary components: the table space and the log. No transaction completes without a log entry. Logs use a tiny fraction of the table space. Logs are response-time sensitive and small in size, so the best practice for a log is a RAID 1/0 with 73 GB 15k rpm drives. Best practice for table spaces is FC 73 GB 15k rpm 4+1 RAID 5.

Ideally a transactional database should spread its data across as many drives as possible. In an OLTP environment the number of drives will have a great effect on response time for any given load. EMC testing uses large drive counts for ideal OLTP environments. Those counts include 120, 240, and 480 drives. The relationship of drive counts to response time is roughly linear. So configurations in our testing have shown that for a given load, half the drives double the response time.

Please check with the database vendor for specific caveats regarding response time and I/O load.

Low to midsize database activity

Databases with user counts up to 160 are considered low to midsize OLTP environments. Generally, the load is a 66 percent to 33 percent mix of reads to writes, with a log write per transaction. Typical I/O counts for a single user would be approximately 30. Logs are response-time sensitive and small in size, so the best practice for a log is a RAID 1 with 73 GB 15k rpm drives. Best practice for table spaces is FC 73 GB 15k rpm 4+1 RAID 5. Please be aware of the response time requirements and size constraints.

Medium database activity

Databases with user counts of 161 to 350 are considered mid-to-large OLTP environments. Generally, the load is a 66 percent to 33 percent mix of reads to writes with a log write per transaction. The typical approximate I/O count for a single user is 30. Logs are response-time sensitive and small in size, so the best practice for a log is a RAID 1 with 73 GB 15k rpm drives. Best practice for table spaces is FC 73 GB 15k rpm 4+1 RAID 5.

High database activity

Databases with user counts of 351+ are considered large OLTP environments. Generally, the load is a 66 percent to 33 percent mix of reads to writes, with a log write per transaction. The typical approximate I/O count for a single user is 30. Logs are response-time sensitive and small in size, so the best practice for a log is a RAID 1/0 with 73 GB 15k rpm drives. Best practice for table spaces for most database environments is FC 73 GB 15k rpm 4+4 RAID 1/0.

Oracle

Consult the *Implementing Oracle on EMC Dell/EMC CX Series Storage Systems* paper for Oracle-specific recommendations.

DB2

Consult *Practitioner's Guide - EMC Database Services DB2 UDB* on EMC Powerlink for DB2-specific best practices.

SQL Server

Consult *EMC Dell/EMC CX Series Database Storage Solutions Microsoft SQL Server 2000 Best Practices* on EMC Powerlink for SQL Server advice.

Exchange

Exchange environments are somewhat OLTP-like, where response time is crucial but different because the I/O pattern is variable in both load and I/O size. Use the exchange best practices

paper published by Microsoft for the correct configuration. You can also use *EMC Dell/EMC CX Series Storage Solutions Microsoft Exchange 2003 Best Practices* on EMC Powerlink.

Backup-to-disk

Backup-to-disk is sequential in nature and typically uses medium (64 KB) to large (256 KB) I/O sizes. For these types of applications use ATA or SATA-II drives, although you can use Fibre Channel drives. To calculate the number of Fibre Channel drives use the “rule-of-thumb” values in *EMC Dell/EMC CX Series Best Practices for Fibre Channel Storage*.

For ATA or SATA-II drives use RAID 3 4+1 or 8+1.

WEB intelligence report repositories

WEB intelligence report repositories are typically data-warehouse-like environments. They collect data for later analysis. The data collection may be either transactional or periodic collection sessions. The analysis is likely to be scans and builds. The I/O profile for these environments is highly sequential reads and writes. During queries, the writes go to a temp area. During data collection, writes go to the main repository.

Use 8+1 RAID 5 LUNs. Fibre Channel drives should be used for high performance, especially if the data-collection method is transactional. If the data-collection method is sequential, then ATA or SATA-II drives should suffice.

Drive count

For RAID 5 the best all-around configuration is five or nine disks. RAID 5 groups can be configured with three to 16 disks.

For RAID 1/0 EMC recommends four to ten disks. RAID 1/0 groups must be grown by adding two disks at a time.

Drive types

For Fibre Channel use 10k rpm drives for back-end loops running at 2 Gb/s or 15k rpm drives for back-end loops running at 4 Gb/s.

Fibre Channel applications include: OLTP, file and print, data warehousing and video.

ATA or SATA-II applications include: backup-to-disk, data warehousing and online archive.

Hot spares

A hot spare is an extra drive for each drive type. EMC strongly recommends a minimum of one hot spare for each type (ATA, SATA-II or FC) of disk in the storage system and one more for every 30 drives of the same type. For hot spares to be effective they must be as large as the largest drive of that type in the storage system. Remember, however, if you have high-speed drives in your configuration, you may also want hot spares that match the speed of the fastest drive. For example, if the system has 146 GB and 300 GB FC drives in the system, the hot spares must be 300 GB to be able to service the 300 GB drives.

Consult the *Dell/EMC CX Series Global Hot Spares – Characteristics and Usage Strategies* white paper for additional details on hot spares.

Drive type

Dell/EMC CX Series supports two different drive types. Within drive types there are different capacities and rotational speeds. This section addresses how these apply to hot spares.

Drive size

EMC has currently certified different FC and ATA/SATA-II drive capacities. EMC does not require any specific disk size within the supported list of disk sizes and types. Select disk sizes and types based on capacity and performance requirements.

Drive technology

The chosen technology must match the performance requirements of the customer application. For OLTP environments, data warehouses, and environments where response time¹ is critical, use FC drives. For backup-to-disk, longer term storage, and applications where response time is not critical, use ATA/SATA-II drives.

Vault and configuration drives

The first five drives in a system are used as database and vault drives. These disks are used for online upgrade operations, replication software management, and Navisphere configuration management and tracking. The RAID group containing the first five drives of the system should not be used for file types other than file-sharing LUNs and should be considered to already have a busy LUN.

Navisphere uses the first three drives for caching online upgrade data. Heavy host I/O during an online upgrade can cause the upgrade to time out, so it is recommended that before an online upgrade commences, the host load be reduced to 100 IOPS per drive. Backups should not be scheduled nor should backups be placed on hold during an online upgrade operation.

Simple and efficient storage system configurations

The minimum CX3-80 configuration consists of one SPS assembly, one SPE with two SPs, and one DAE with five drives.

For DAEs with hot spares, configure 1x 8+1 and 1x 4+1 and one hot spare. For DAEs without hot spares, create 3x 4+1 RAID groups.

¹ Response time is defined as the difference in time from when the user begins the computer request until the user gets the data requested or a write completion notification. Response time is how a human measures performance.

Alternating 8+1 and 4+1 Raid groups

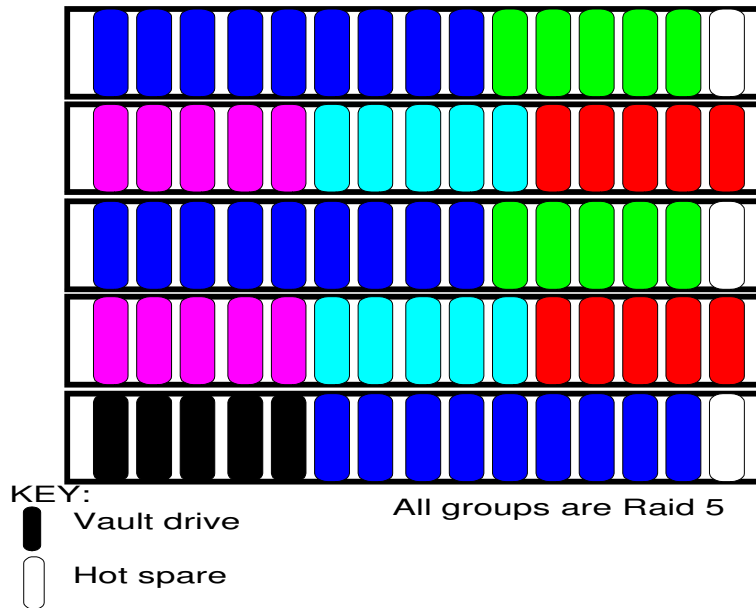


Figure 1. Alternating 8+1 and 4+1 RAID groups

8+1 Raid groups

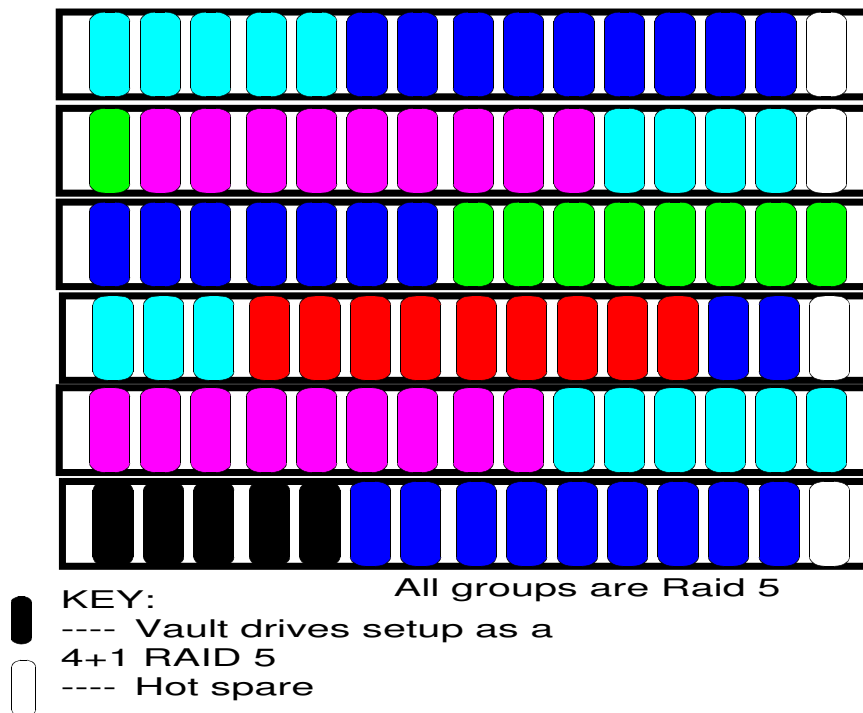


Figure 2. 8+1 RAID groups

Raid 1/0 configuration

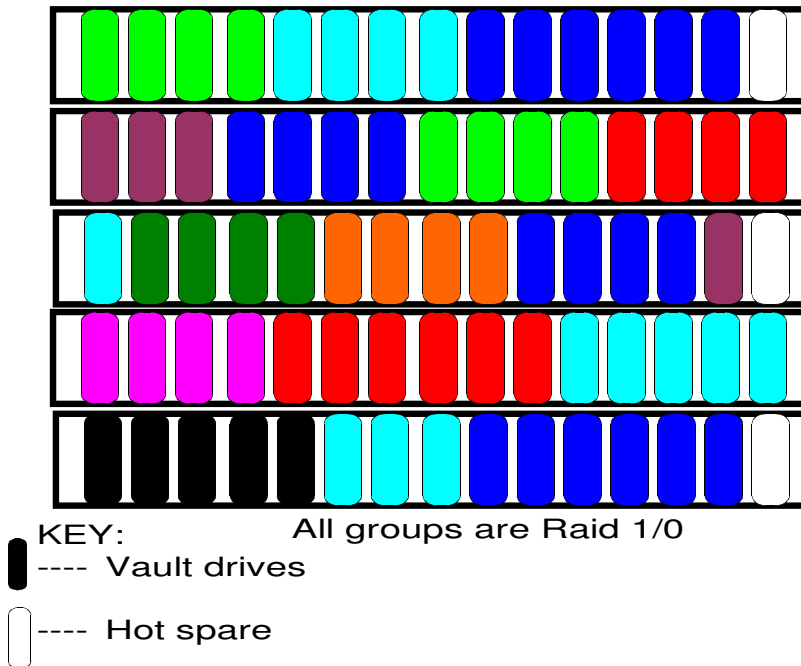


Figure 3. RAID 1/0 configuration

Storage Groups

A single-host system connected to the CX3 Series storage system will be added to a dedicated Storage Group. Multi-system cluster nodes will be added to a shared resource Storage Group for disk access. EMC does not provide a Storage Group naming standard as a best practice.

Create a Storage Group

Storage Groups must be defined before you can assign LUNs to a specific host. Table 4 describes the user-defined and default parameters to create a Storage Group.

Table 4. Create a Storage Group

	Set	Default setting	Notes
Create Storage Group			
Storage Group name	See note	Storage Group 1	Use customer naming standard for Storage Groups.

Storage Group settings

Table 5 describes the user-defined and default parameters for a Storage Group.

Table 5. Storage Group settings

	Set	Default setting	Notes
General Tab			
Storage Group Name	See note	Storage Group 1	Use customer-naming rules, if any exist.
LUNs Tab			
Show LUNs	Default	Not in other Storage Groups	Exception to this is use of clusters. Refer to the "Clustering" section.
Available LUNs	Add at least one LUN.	Nothing selected	Select all LUNs the host needs to see.
Hosts Tab			
Show Hosts	Default	Not connected	
Select Hosts	Add at least one host.	Nothing selected	A host can only be used in one Storage Group.

Caching

For 80 percent of the configurations, enable both read and write cache. Assign as much as possible to write cache. Split the remaining memory between SP A read cache and SP B read cache. Ensure that read and write cache is enabled for each LUN.

Cache settings and other SP information

Table 6 describes the user-defined and default storage system properties.

Table 6. Cache settings and SP information

Properties	Set	Default setting	Notes
General Tab			
Name	See note	CXtype_ipSPA_ipSPB	Use Customer Naming rules
Statistics Logging	Box checked	Box not checked	
Cache Tab			
Page Size	Default	8 KB	
Low Watermark	Default	60	
High Watermark	Default	80	
Enable Watermarks	Default	Box checked	
SP A Read Cache	Box checked	Box not checked	
SP B Read Cache	Box checked	Box not checked	
Write Cache	Box checked	Box not checked	
HA Cache Vault	Default	Box checked	
Memory Tab			
SP A/SP B Read Cache	See Note	0	Leftover memory from write cache should be split between SP A and SP B read cache.
Write Cache	As much as allowed	0	Depends on Dell/EMC CX Series model
RAID 3 Memory	Do not use	0	
Storage Access Tab			
Access Control Enabled	Box checked	Box not checked	
Software Tab			
Commit	See note		This button pertains to upgrades. This button is grayed out if commit has already been done.
Revert	See note		This button pertains to upgrades. This is grayed out if commit already done.

Software online upgrades

EMC recommends Host Agents be upgraded as soon as possible so they are the same revision as the Navisphere revision on the storage system. The management station must be running the same revision as the highest revision of Navisphere running in the domain. Navisphere software is backwards-compatible, so the management station may be upgraded before the storage systems are upgraded.

Upgrades to the storage system can be achieved by an operation called **online upgrade**. Online upgrades require that no more than 100 I/Os are running on the first five disks in the storage system. EMC recommends that all online upgrades be performed at off-peak hours.

The Dell/EMC CX Series allows reverting to a prior revision as long as the FLARE code has not been committed. The tradeoffs are:

- No changes are allowed to the Dell/EMC CX Series.
- New features in the new revision of FLARE may not be usable.

Not committing allows you to confirm that the basic I/O functions work in your environment. Once committed, the FLARE code can only be upgraded. EMC recommends committing the new FLARE revision at installation. Therefore, the ability to revert should be considered the exception.

Ensure all host, layered application and switch components (driver, agent, and firmware) are upgraded to maintain compatibility.

MONITORING

This section discusses the different management methods available to the user.

Management system

To manage the storage system, use the installed management software.

Access to the management system requires browser software. EMC recommends Internet Explorer, Netscape, and/or Mozilla, depending upon the OS.

Use Java JRE revision 1.4.2 for the browser. Note that this may change faster than this document is revised. Navisphere will tell you if you are running the wrong revision of Java and point you to the web site for the correct version. If you are not running the correct revision of Java, the Navisphere UI will send the user to the Java site. It is therefore also recommended that the system running the browser have access to a public network.

A Dell/EMC CX Series agent is not required on the system using a browser connection to manage a Dell/EMC CX Series domain unless the system has a FC or iSCSI connection to the Dell/EMC CX Series and is used to manage a legacy Dell/EMC CX Series via portals.

Off-array management server

EMC supports management of a Dell/EMC CX Series from either a management station or the storage system itself. The following are the EMC recommendations for off-array management.

- To manage a Dell/EMC CX Series using the off-array management server, that station must be a member of a Dell/EMC CX Series domain.
- For management stations running pre-release 19 software, EMC recommends a single management server manage no more than 30 storage systems at one time. The storage systems can be managed as one session if they are all placed in a domain.
- For management stations running post-release 19 software, EMC recommends no more than 72 total storage systems in the multi-domain configuration. The storage systems may be in one domain or several domains.
- The server should be a Windows server running at least Windows 2000 SP4. The server should have at least a 1-GHz CPU and 512 MB of RAM.
- The server used for off-array management should not be running any production application. You can use the same server as the CLARAlert® workstation. Dell/EMC CX Series should be in the same domain as the management station of Dell/EMC CX Series.

In-array management

EMC supports management of the Dell/EMC CX Series directly. The storage system limits recommended for off-array management also apply to in-array management. EMC does not have a best practice for which management method to use; both off-array and in-array are effective ways to manage a Dell/EMC CX Series. Follow your corporate standards for Navisphere user account names. This will cut down on confusion and allow for better accounting and security measures.

Host management

The Navisphere Management Suite provides several options for host management based on the level of service desired, the mode in which the host component will run, and the level of required host-storage system communication. EMC recommends the use of the Navisphere Host Agent as this is the most full-featured option.

The Host Agent automates the establishment of connectivity with the storage system by identifying the host. The agent also provides mount information to the LUNs. This information is used by the storage manager to map storage system LUNs to the host file systems. EMC recommends that the Host Agent software and the Navisphere software running in the Dell/EMC CX Series be at the same revision.

The Host Agent runs as a daemon/service on the host system. If a persistent service is not desired, EMC recommends the Server Registration Utility for host management. This utility is run on demand and performs host registration and LUN/volume mapping functions.

Host Agent settings

Set the default privilege settings as shown below:

```
system@SPA_IPAddress
system@SPB_IPAddress
system@CallHomePC_IPAddress
WindowsUser@CallHomePC_IPAddress
```

Set host agents polling to 120.

Management scope

The scope provides access to either local or full domain access.

- **Global** – User account can access all subsystems in the domains.
- **Local** – User can only access the particular subsystem.

Use Global scope for work in any domain.

Local diagnosis EMCRemote

EMCRemote is a remote access utility used by EMC support for dial-in diagnostics and troubleshooting. Originally developed for the Symmetrix® line, it is now embedded into the Dell/EMC CX Series CLARAlert package.

Event notification

By default, notification is not enabled. Have EMC install CLARAlert software on the management workstation. The Phone Home software provides a template of events that EMC considers Phone Home worthy events. Do not add to this list.

Event Monitor

The default template provides logging capability only. Critical and hard errors will be logged. Do not change this template.

Event Monitor settings

In the **Monitors** tab of the Navisphere UI, right-click **Templates** and select **Create New Template**. ClarAlert will use Event Monitor, so at the minimum use these settings. If you wish to extend alerting internally to your company, use the other options shown in Table 7.

Table 7. Event Monitor settings

	Set	Default setting	Notes
General Tab			
Events	Explicit Events	General	Events selected by EMC
Log to System Logs	Default	Box unchecked	
Combine Events for Storage System.	Box Checked	Box unchecked	
E-mail Tab			As needed in special cases, consult documentation.
Email To	Default	Field empty	
Cc	Default	Field empty	
From	Default	Field empty	
Subject	Default	Field empty	
SNMP Server	Default	Field empty	
Use Default Message	Default	Box checked	
Paging Tab			As needed in special cases, consult documentation.
Use Modem	Default	Box checked	
All Other Fields	Default	Field empty	
SNMP Tab			As needed in special cases, consult documentation.
All Fields	Default	Field empty	

Phone Home

Phone Home software is provided by EMC and installed on a Windows server during the Dell/EMC installation at a new site. The newly installed Dell/EMC must be setup to talk to the Phone Home system during installation. The Phone Home function means the environment notifies EMC of significant events on the Dell/EMC. These events are preset during Phone Home installation but the customer may add more events if they desire.

Best practice: Do not change the event list.

Phone Home settings

Table 8 describes the user-defined and default Phone Home settings.

Table 8. Phone Home settings

	Set	Default setting	Notes
Events	Explicit Events	General	Events were selected by EMC
Log to System Logs	Do not change	Box checked	
Combine Events for Storage System.	Box Checked	Box not checked	
E-mail Tab			
Email To	Do not change	Field empty	
Cc	Do not change	Field empty	
From	Do not change	Field empty	
Subject	Do not change	Field empty	
SMTP Server	Do not change	Field empty	
Use Default Message	Do not change	Box checked	
Paging Tab			
Use Modem	Do not change	Box checked	
All Other Fields	Do not change	Field empty	
SNMP Tab			
All Fields	Do not change	Field empty	

SNMP management

Dell/EMC CX Series supports SNMP read requests and SNMP trap notifications.

To use SNMP to query a Dell/EMC CX Series, enable the Fibre Alliance MIB package. This package is pre-installed on all currently shipping Dell/EMC CX Series storage systems. Enabling the Fibre Alliance MIB is not required for SNMP traps. See Table 9 for the Fibre Alliance MIB settings.

Table 9. SP properties: Network tab

	Set	Default setting	Notes
Network Tab			
SP Network Name	Do not change	SP Serial Number	This is set at installation
IP Address	Do not change	192.168.1.1	This is set at installation
Gateway	Do not change	Field empty	This is set at installation
Subnet Mask	Do not change	Field empty	This is set at installation
Enable/Disable processing of SNMP MIB read requests	Box checked	Box unchecked	

For SNMP notification, enable SNMP traps in event monitor. Use the Phone Home template for the list of errors. See Table 10 for the SNMP settings.

Table 10. SNMP settings

	Set	Default setting	Notes
SNMP Tab			
SNMP Management Host	Enter target host	Field empty	
Community	Enter community	Field empty	

SNMP traps

Object Identifiers (OIDs) uniquely identify any object on a network. See Table 11 for a list of OIDs contained in a Dell/EMC CX Series trap.

Table 11. OIDs contained in a Dell/EMC CX Series trap

OID	Description
.1.3.6.1.4.1.1981.1.4.3	Host Name of Agent Sending the Trap
.1.3.6.1.4.1.1981.1.4.4	SCSI Device Name ("K10" in FC4700 Arrays)
.1.3.6.1.4.1.1981.1.4.5	Event Code Number
.1.3.6.1.4.1.1981.1.4.6	Event Code Text
.1.3.6.1.4.1.1981.1.4.7	Array Reporting the Event

All Dell/EMC CX Series running release 19 firmware contain five trap objects. Each trap definition is identical as shown in Table 11; however the OID for the trap itself is based on the event severity. See Table 12 for a list of OIDs for all Dell/EMC CX Series trap objects.

Table 12. OIDs for Dell/EMC CX Series trap objects

OID	Description
.1.3.6.1.4.1.1981.0.2	All events for pre-release 19 storage systems and events when severity cannot be determined for release 19 and higher storage systems.
.1.3.6.1.4.1.1981.0.3	Informational
.1.3.6.1.4.1.1981.0.4	Warning
.1.3.6.1.4.1.1981.0.5	Error
.1.3.6.1.4.1.1981.0.6	Critical Error

Management notification

The default for notification is **none**. Install Phone Home or setup notification for the system.

Sessions limit

No more than 10 simultaneous privileged users should be running Navisphere in a domain at the same time, and no more than five to any one Dell/EMC CX Series. Only one user should modify the configuration at any one time. If two privileged users try to modify the configuration at the same time, Navisphere will not prevent it.

SECURITY

This section discusses access rights and protection mechanisms in the Dell/EMC CX Series.

Storage systems

All Dell/EMC CX Series should be in a secure room. The only “back door” access to the Dell/EMC CX Series requires physical access to the serial port.

Enable security on each installed Dell/EMC CX Series. This establishes the username-password access to the user interface (UI).

Management

The Dell/EMC CX Series provides secure management by password protection, layered access rights, and SSL encryption of all management-related communications. For example, to access the storage system the user must type a username and password. Both of these will be encrypted prior to passing the access data over the IP connection.

Domains

Each Dell/EMC CX Series must be in a storage domain.

Access levels

There are multiple access levels for users in Navisphere.

- **Administrator**—All manager privileges. The administrator is allowed to add and delete users and set privileges.
- **Manager**—The manager is allowed to change any setting in the Navisphere UI or CLI except user-related settings.
- **Monitor**—The monitor can see the Dell/EMC CX Series in the domain but cannot change any attribute.

Only one or two people should have administrator privileges. Production support staff should have manager privileges. All others should have either monitor privilege or be undefined.

Audit logging

Navisphere has enhanced logging and auditing capabilities; the log entries are sent to the Navisphere Management Server event log.

Password control

Passwords can be any alpha/numeric symbol. They are case sensitive. EMC has no preference for password naming. Passwords should be changed at regular intervals and after termination of a privileged user.

Access lost or forgotten passwords

If passwords for all privileged users are lost, the only way to access the storage system is via the SP's serial port.

CONNECTIVITY

This section discusses the different connectivity methods.

Management

EMC has customers that use either the on-array or off-array Navisphere management. Provide one LAN connection to each Dell/EMC CX Series SP. Use Host Agents to automate connectivity of hosts at boot time. The agent is not always necessary if a connection was already established, but it will be needed after a power fail and a restart of the Dell/EMC CX Series. Therefore, it is recommended to always run the Host Agent.

Protocols

There are two protocols available for Dell/EMC CX Series storage systems: Fibre Channel (FC) and iSCSI. The FC protocol is available for all CX3 Series systems. Limited CX3 devices allow iSCSI as an option. Refer to the product listing on <http://www.EMC.com> for a complete list of CX3 Series systems that support iSCSI.

iSCSI connections

iSCSI connections from the storage system can be made to either network interface cards (NIC) or TCP/IP offload engines (TOE). Use cat 6 cables to connect to the Dell/EMC CX Series iSCSI ports.

The Microsoft, QLogic, and Dell/EMC CX Series software all provide ping and trace-route capability to assist with troubleshooting connections. Use these tools to check connectivity if the host or the Dell/EMC CX Series report that they do not have a connection.

NIC cards

The NIC port speed can be set to 10/100/1000 Gb with CX3 iSCSI/FC storage systems. Use IEEE 802.3 certified NIC cards.

TOE cards

Check the *EMC Support Matrix* for the current list of supported TOE cards. When using a TOE card, you need to remove the native (Microsoft) iSCSI driver from the host. Booting from SAN requires a TOE card.

iSCSI security

Currently the only iSCSI security supported is the Challenge Handshake Authentication Protocol (CHAP). CHAP requires a secret password on both the host and the Dell/EMC CX Series be set during installation. At a minimum, use initiator-based CHAP security. The secret password must be 12-to-16 characters long and must be correctly typed on both sides of the connection. Typos in the secret name will cause the connection to break when CHAP is enabled.

Fibre Channel connections

EMC supports Emulex and QLogic host bus adapters (HBAs). For specific types, firmware revisions, and driver revisions, refer to the *EMC Support Matrix*.

Use current drivers and firmware. Current driver and firmware revs for HBAs can be found in the *EMC Support Matrix*.

For HBA device driver settings, driver levels, or components in the data path, tape or disk configurations should include any patches, firmware updates, or configuration settings for each component in or on the data path.

Emulex HBAs

Use the EMC setting in the HBA configuration utility. Supported Emulex HBAs and driver revs can be found in the *EMC Support Matrix* on EMC Powerlink.

Settings for Emulex

Use the settings defined in the Emulex driver installation manual for the specific HBA.

Variations for tape

When using tape, there are two more settings needed:

Use PLOGI instead of PDISC after LIP – Always set this when connecting tape devices to the HBA.

Disable target reset for tape devices – This should be set when using tape devices and VERITAS cluster software. If cluster software is not used, this parameter should not be set.

QLogic HBAs: QLogic settings

Some HBAs may not come with the EMC settings. If the settings on the HBA match the EMC settings, then the HBA was preconfigured to the correct settings for connection to a Dell/EMC CX Series.

Data path

The data path should be redundant to reduce the chance for failures. Multipath software runs on the attached host. A minimum of two paths—one to each SP—must be created to provide redundancy. For a complete list with software revisions refer to the *EMC Customer Matrix*.

EMC recommends using PowerPath® for multipath host software for the supported hosts. To take advantage of PowerPath load balancing and redundancy, create four paths to the Dell/EMC CX Series. This can be accomplished with two HBAs and two ports per SP.

EMC recommends dual-attached hosts.

Host bus adapters

EMC also recommends persistent binding from the host to the Dell/EMC CX Series.

Recommended HBAs

The most prevalent HBAs in use with and supported by EMC are QLogic and Emulex. Other native adapters include the Sun adapters (which are QLogic HBAs), and IBM adapters, (which are Emulex HBAs).

Refer to the *EMC Support Matrix* for specific HBA and driver support for the hosts in your environment.

Dell/EMC CX Series SP ports

Use the lower numbered ports for general connectivity.

The highest numbered port is used for MirrorView connections. All ports can be used for other general connections but it is recommended to keep the high port exclusively for mirroring. If MirrorView is not installed then the high ports should be used for general connectivity.

Fibre Channel SAN link speed

The recommended Dell/EMC CX Series link speed is 4 Gb/s. In any case, EMC recommends setting the link speed to **auto**. If the switch does not support 4 Gb/s, then the Dell/EMC CX Series SP port will automatically set itself to the lower speed. Currently, the port-speed choices are auto, 4 Gb/s, 2 Gb/s and 1Gb/s. Set the switch to **auto negotiate**.

Zoning

EMC strongly recommends World Wide Name (WWN) zoning on all switches in the SANs. EMC also recommends single initiator zoning.

Port-to-HBA relationships

EMC recommends that each HBA be connected into a different fabric. Connect each SP into both fabrics. This provides a total of four connections to the Dell/EMC CX Series for each LUN. For a CX3-80 this kind of connectivity allows 256 highly available hosts attached to the Dell/EMC CX Series. Additional connections may be added for performance, but this will decrease the number of hosts attached to the Dell/EMC CX Series. The minimum configuration is one HBA path to each Dell/EMC CX Series SP for a minimum of two paths to the Dell/EMC CX Series.

Four Dell/EMC CX Series ports should be seen by each host; two ports from SP A and two ports from SP B.

Fabric connectivity

EMC allows a maximum of three hops between the Dell/EMC CX Series and the host. Minimize the number of hops whenever possible.

Multipath software

EMC supports several multipath software products. EMC recommends PowerPath for all Dell/EMC CX Series supported hosts as a best practice.

PowerPath

PowerPath provides failover and load-balancing capability across all the ports configured to the Dell/EMC CX Series. PowerPath supports a maximum of eight active and eight passive paths to any Dell/EMC CX Series LUN.

CLAR_Opt is the default load-balancing configuration. Do not change this setting.

Figure 4 shows the recommended path layout for a single host. This configuration provides both failover and load-balancing capability.

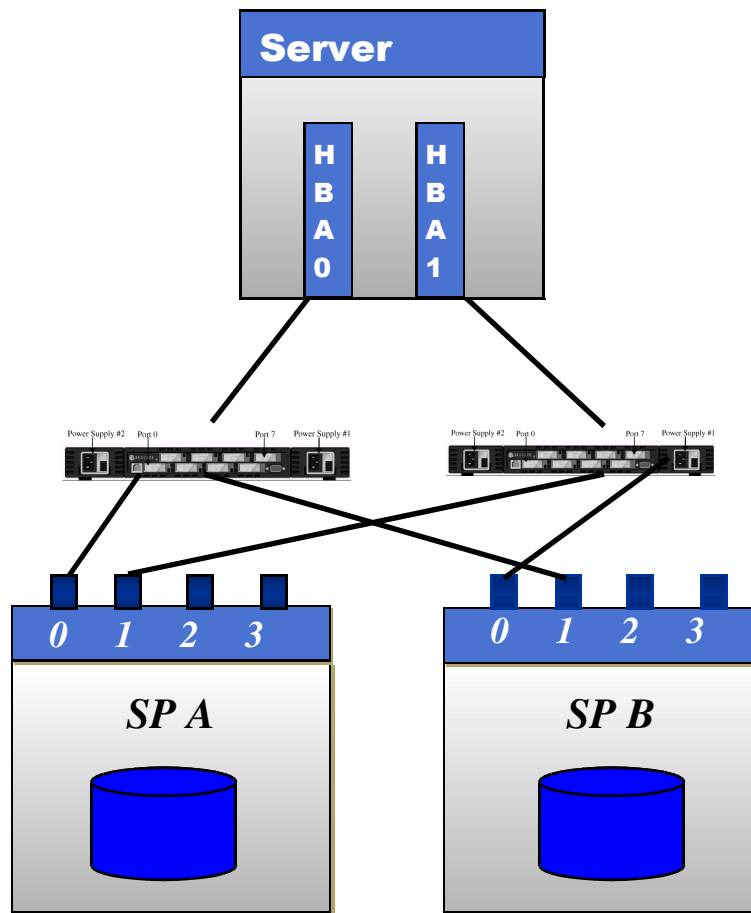


Figure 4. Single host path layout

DMP

VERITAS DMP is supported on Windows, HP/UX, Solaris and Linux systems. DMP can coexist with PowerPath, but no load balancing will be provided by DMP. To use DMP to a Dell/EMC CX Series without PowerPath, a library must be loaded from the VERITAS website. The file is called: CLR-ASL. The same file is found on the platform CD.

LUN presentation

By default, LUNs will be numbered on the host side in order of discovery. This means that in most cases, they will be discovered in display order of the Storage Group. In environments that change, the host LUN ID may change as LUNs are added and removed from a Storage Group. Best practice is to keep the Storage Group as static as possible. This can be accomplished with thorough design of the RAID group and LUN layout for each host brought into the environment.

The host LUN ID can also be forced with the **-HLU** switch in the NAVICLI or with Navisphere Manager; however, this is the exception.

RAID group to back-end bus connections

Create the RAID groups within the same DAE. This will force all the members of a RAID group to be on the same bus. Putting all the disks on the same bus simplifies the configuration. There is

little performance advantage to splitting the disks across all the buses, so configuring the disks across all different buses is not worth the added complexity.

Back-end connectivity

Back-end connectivity is controlled by the EMC installer to help minimize complexity and maximize balance of the number of disks on a pair of back-end buses.

Let Dell/EMC CX Series manage its own back end connectivity.

Clustering

EMC supports several clustering products. Fundamentally, cluster support can be accomplished on a Dell/EMC CX Series by sharing common LUNs in the Storage Group. If the members of the cluster need their own separate storage, separate Storage Groups must be used and the common LUN is put into both Storage Groups. A host can only be in one Storage Group per storage system, but LUNs can be in multiple Storage Groups. When LUNs are shared between the hosts, it is up to the sharing hosts to manage the data integrity of shared LUNs. Therefore create two Storage Groups—one for each host—and put the shared LUNs in both Storage Groups. When selecting the LUNs for the second Storage Group, select **all** in the **show LUNs** box and select the shared LUNs. When you select **all** a warning informs you that “multiple hosts” will be able to see the LUNs. Click **OK**. Be sure that the LUNs selected in the Navisphere UI for the second host’s Storage Group are selected in the same order as the selections in the first host’s Storage Group. Selecting the LUNs in a different order may cause problems with the cluster software.

Choose a Storage Group name that reflects the host and cluster name. For example, you might call the Storage Group containing the first host in a Sun cluster: Host1-of-Suncluster.

The following sections show the supported cluster software. Cluster software is host or OS based. Typically cluster software uses local copies of the data and does not participate in remote copies of the data. Clusters with Dell/EMC CX Series do not typically use Dell/EMC CX Series replication software for heartbeats. Shared disks can be used as heartbeat or mailbox storage, but that does not depend on any Dell/EMC CX Series-specific capability other than shared drives. Heartbeats/mailboxes for the cluster must be shared drives.

Sun clusters

Do not use Sun Cluster if Solaris is booting from the Dell/EMC CX Series.

You must use FC-SW to support OPS, RAC, or greater than two-node clusters.

Some Solaris 9 clustered environments require patches. Refer to the *EMC Support Matrix* for patch listings for particular hosts and configurations.

VERITAS clusters

Do not use VERITAS Cluster if Solaris is booting from the Dell/EMC CX Series.

IBM clusters

Two IBM Cluster packages are available to users: HACMP and GPFS.

HACMP

Booting from the Dell/EMC CX Series is allowed in HACMP clusters. These boot drives should be only available to that host.

GPFS

GPFS configurations usually span multiple hosts and storage systems to create the parallel file systems structure. These configurations are usually in the 40–100 TB size range. Multiple

Dell/EMC CX Series can be used to accommodate these configurations. This is considered outside the 80 percent rule for best practices. Consultation is required for these configurations.

Windows clusters: Microsoft Cluster Server

Current SCSI port or SCSI miniport drivers cannot share the same path for the boot device as for the shared cluster devices. As a result, a dedicated adapter is required for each path (if booting from internal disk, it is already on a separate Adapter). For example, each host has its own boot LUN and additional data LUNs that they share. A non-HA configuration requires two HBAs for each server: one for a dedicated path to the boot device and another for a path to the data devices. An HA configuration for the data LUNs requires three HBAs each (two data HBA paths). An HA configuration to include the boot LUNs requires four HBAs on each server. Each Server needs separate paths: Two HBA paths for the data and two paths for the boot LUN.

Storport can share the same path for the boot device and the shared cluster devices.

Microsoft Windows Server 2000/2003 clusters using Microsoft Cluster Services (MSCS) cannot make use of dynamic disks. Use metaLUNs and Microsoft's diskpart utility to provide expandable, striped, and concatenated volumes for these systems.

Linux (RedHat) clusters

Use Enterprise Linux Cluster.

HP TruCluster

Dell/EMC CX Series requires HP TruCluster software.

SCALING: STORAGE SYSTEM SIZES

Dell/EMC CX Series storage systems can be defined many ways. Earlier sections of this white paper defined usage profiles. This section defines three systems based on size or capacity. The systems are: small, medium and large.

Small storage system

A small storage system consists of one to three DAEs. The first DAE will have Fibre Channel drives. The other two can have either ATA/SATA-II or FC drive types.

Medium storage system

A medium storage system ranges in size from four to 11 DAEs, to the maximum being a full single rack.

Large storage system

A large system uses two or three racks, which are added footprints. The large system contains 12 to 32 DAEs. The CX3-80 can support up to 480 drives per storage system.

Host connectivity

Environment sizes are defined as follows:

- Small environments have 1 to 10 hosts.
- Medium environments have 11 to 30 hosts.
- Large environment have greater than 31 hosts.

Exchange environments

Exchange environments are defined by size and load. Size is measured in numbers of users, and load is defined as the number of I/O per second per user.

- Small environments are less than 1,000 users.
- Medium environments have 1,000 to 5,000 users.
- Large environment have greater than 5,000 users.
- Light loads are less than .18 IOPS (OWA and POP3) per user.
- Medium loads consist of .4 to 1 IOPS (Outlook Rich Client) per user.
- Heavy loads consist of greater than .75 IOPS (Outlook Rich Client) per user.

Table 13 provides the recommended storage system configuration based on the Exchange environment and workload. The configurations are defined in the section “S.”

Table 13. Storage system configuration for various Exchange environments and workloads

	Small environments	Medium environments	Large environments
Light loads	Small configuration	Medium configuration	Large configuration
Medium loads	Medium configuration	Large configuration	Large configuration
Heavy loads	Tuning required*	Tuning required*	Tuning required*

* These are exceptions.

Consult the *EMC Dell/EMC CX Series Storage Solutions Microsoft Exchange 2003 Best Practices* white paper for additional details.

DISK-ARRAY ENCLOSURE (DISK SHELVES)

Availability

Scaling is accomplished on the Dell/EMC CX Series by adding DAEs. DAEs can be added while the system is running, but if possible, EMC recommends shutting down the attached hosts to the system when making hardware changes to the environment. Hardware changes should be coordinated with EMC service.

Drive types per DAE

Only one type of drive— either Fibre Channel, SATA-II or ATA—can reside in a specific DAE. The first DAE must have *only* FC drives.

Drives per DAE

A DAE supports up to 15 drives. EMC does not require any drives in a DAE, other than in the first DAE. EMC recommends a minimum of four drives per DAE.

Drive types per storage system

EMC recommends matching the drive type to the type of application environment. The first tray must contain Fibre Channel drives.

Drives per RAID group

Select 5, 9, 2, or 1 for the number of drives per RAID group.

- 5 = 4+1 RAID 5
- 9 = 8+1 RAID 5
- 2 = 1+1 RAID 1 or RAID 1/0
- 1 = Hot spare

Cabinets

A second cabinet is required once a total of 11 DAEs has been reached.

A third cabinet is required once a total of 13 DAEs has been reached in the second cabinet.

CONCLUSION

This white paper describes the Dell/EMC CX Series storage system and the connection settings that satisfy 80 percent of the Dell/EMC CX Series environments. This is a recommendations or best practices document— not a requirements document. It is also a living document, which means that as better settings are discovered, changes will be made to the document.

REFERENCES

Replication software

Several replication software applications run on Dell/EMC CX Series. The following list of best practices papers address the special capabilities of replication software. Consult the appropriate paper for the replication software being used.

Replication Manager and SnapView for Replication on Dell/EMC CX Series Storage Arrays - Best Practices Planning

Dell/EMC CX Series Reserved LUN Pool Configuration Considerations for SnapView, SAN Copy, and MirrorView/Asynchronous

EMC Dell/EMC CX Series SnapView Clones - A Detailed Review

EMC Dell/EMC CX Series SAN Copy - A Detailed Review

EMC Dell/EMC CX Series MirrorView/Asynchronous Disaster Recovery Software

MirrorView Knowledgebook: CX Series FLARE 19 and CX3 Series FLARE 22 - Applied Technology

Other references

EMC Dell/EMC CX Series Best Practices for Fibre Channel Storage on EMC Powerlink

EMC Dell/EMC CX Series Storage Solutions Microsoft Exchange 2003 Best Practices on EMC Powerlink

EMC Dell/EMC CX Series Data Replication Scenarios for Oracle Deployments - Technology Concepts and Business Considerations on EMC Powerlink

SQL Server Data Warehouse Deployments with EMC Dell/EMC CX Series Storage Systems Applied Technology White Paper on EMC Powerlink

The information contained in this document, including all instructions, cautions, and regulatory approvals and certifications, is provided by EMC Corporation and has not been independently verified or tested by Dell. Dell cannot be responsible for damage caused as a result of either following or failing to follow these instructions. All statements or claims regarding the properties, capabilities, speeds or qualifications of the part referenced in this document are made by EMC and not by Dell. Dell specifically disclaims knowledge of the accuracy, completeness or substantiation for any such statements. All questions or comments relating to such statements or claims should be directed to EMC Corporation.